

RECEIVED
CENTRAL FAX CENTER

DEC 05 2006

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICANT: Matt Ayers ART UNIT: 2144
SERIAL NO.: 09/575,839 EXAMINER: Thanh Nguyen
FILED: 05/22/2000
FOR: Method and System for Directing
Requests for Content to a
Content Server Based on
Network Performance

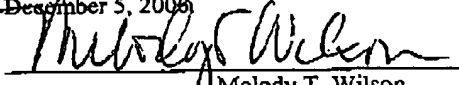
ATTORNEY DOCKET NO.: 52224/296056

Commissioner for Patents
Mail Stop Appeal Brief
P.O. Box 1450
Alexandria, VA 22313-1450

CERTIFICATE OF TRANSMITTAL

I hereby certify that this correspondence is being facsimile
transmitted to the U.S. Patent and Trademark Office's
Centralized Facsimile Number of 571.273.8300 on

December 5, 2006


Melody T. Wilson

AMENDED APPEAL BRIEF

Sir:

This amended Appeal Brief is submitted in response to the Notice of Non-Compliant Appeal Brief mailed on October 26, 2006, which required additional explanation of the subject matter defined in each of the independent claims involved in the appeal. This brief amends the section entitled "Summary of Claimed Subject Matter." A request for a one-month extension of time accompanies this response.

RECEIVED
CENTRAL FAX CENTER

DEC 05 2006

Table of Contents

	<u>Page</u>
I. Real Party In Interest.....	3
II. Related Appeals and Interferences.....	4
III. Status of Claims.	5
IV. Status of Amendments.	6
V. Summary of Claimed Subject Matter.....	7
VI. Grounds of Rejection to be Reviewed on Appeal.....	11
VII. Argument.....	12
A. Rejection under 35 USC § 103(a) over Colby, in view of Spicer.....	12
B. Rejection under 35 USC § 103 (a) over Colby, Spicer and Jang.....	17
VIII. Claims Appendix.....	20
IX. Evidence Appendix.....	36
X. Related Proceedings Appendix.....	37

RECEIVED
CENTRAL FAX CENTER

DEC 05 2006

I. Real Party In Interest

The real party in interest is Internet Network Services Corporation, the successor in interest to Internet Network Services, the assignee of record.

RECEIVED
CENTRAL FAX CENTER
DEC 05 2006

II. Related Appeals and Interferences.

There are no related appeals or interferences.

RECEIVED
CENTRAL FAX CENTER

DEC 05 2006

III. Status of Claims.

Claims 1-21, 32-53 and 64 are pending and Claims 22-31 and 54-63 have been cancelled. Claims 1-21, 32-53 and 64 are rejected.

RECEIVED
CENTRAL FAX CENTER
DEC 05 2006

IV. Status of Amendments.

No amendments are outstanding.

RECEIVED
CENTRAL FAX CENTER

DEC 05 2006

V. Summary of Claimed Subject Matter.

The present invention is directed to speeding the delivery of content to users over a network. Many prior art solutions select a server that is geographically closest to the user in an attempt to speed delivery. However, this approach does not always provide the fastest response time. The present invention selects a content server based on criteria other than which server is geographically or topologically closest to the user. Page 3, lines 16-26.

In one embodiment a client desires to access content associated with a customer's web site. A client is a computer that accesses as a network, such as the Internet, and a customer is a network server, such as a web site server. Page 6, lines 6-10. Every time a client accesses a content server the client provides performance data about that content server that can be observed by a spanner. The performance data is used to estimate performance for other clients within a cluster or set of IP addresses that contain the client's IP address. Page 11, lines 11-15.

It is not practical to monitor the performance from every IP address or every set of addresses (*e.g.*, prefixes) due to the large number of measurements that would be required. However, if prefixes are close together geographically, then an inference can be made that physically proximate prefixes will experience similar network performance. Page 11, lines 21-29. The claimed invention infers performance for sets of IP addresses that are physically close, but for which no actual measurement have been made. Page 13, lines 3-5.

Figure 2 illustrates a situation where measurements are available for certain sets of IP addresses, P1 and P2, but not for all sets, such as P3. However, since P3 is within the circle of inference for P1 and P2, the measurements for P1 and P2 can be used to identify a server for P3 based on measurements related to P1 and P2. Page 13, lines 6-10.

Claim 1 is set forth below and requires directing a first network client to one of the content servers if cost measurements are available that measure operational characteristics of the network based on communication between the first network client and one or more of the content servers, otherwise directing the first network client to one of the content servers using one or more cost measurements that measure operational characteristics of the network based on communication between a second client one or more of the content servers, where the second client is physically proximate to the first network client and. Claims 12, 33 and 44 contain similar recitations to Claims 1 and are set forth in their entirety in the Claims Appendix. The invention recited by Claims 1, 12, 33 and 44 is described in the specification, *inter alia*, on page 11, line 27 – page 13, line 10 and Figure 2.

1. A method for directing a first network client requesting access to content to one of a plurality of content servers that can provide said content, comprising:

if one or more cost measurements are available that measure operational characteristics of the network based on communication between the first network client and one or more of the plurality of content servers, then directing the first network client to a said one of said content servers based on the one or more cost measurements;

otherwise, directing the first network client to a said one of said content servers using one or more cost measurements that measure operational characteristics of the network based on communication between a second client that is physically proximate to the first network client and one or more of the plurality of content servers.

Figure 5 illustrates an “intelligent inference” method where latency measurements are available between content server, Ca, and certain sets of IP addresses, S1, S2, S3. The latency measurements and distances between the sets of IP addresses, S1, S2, S3, and another set of IP addresses, S4, are used to infer a latency measurement between the content server, Ca, and the set of IP addresses, S4. Page 15, line 18-line 29.

Claim 32 is set forth below and requires inferring operational characteristics associated with a plurality of network clients to an inferable network client by determining physical distances between said network clients and the inferable network client and computing a weighted average of said latency measurements as a function of said distances, wherein said weighted average comprises an estimate of the latency between said network server and said inferable network client. Claim 64 contains similar recitations to Claim 32 and is set forth in its entirety in the Claims appendix. The invention recited by Claims 32 and 64 is described in the specification, *inter alia*, on page 15, line 18 – page 16, line 15 and Figure 5.¹

¹ Note that the references to Fig. 10 and Fig. 4 in the paragraph beginning on page 15, line 18 of the originally filed specification should be references to Fig. 5.

32. (Previously Presented) A method for inferring operational characteristics associated with a plurality of network clients to an inferable network client, comprising:

(a) measuring network latency between a network server and the plurality of network clients;

(b) determining physical distances between said network clients and the inferable network client; and

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said network server and said inferable network client.

VI. Grounds of Rejection to be Reviewed on Appeal.

- Whether Claims 1-8, 10-18, 20, 21, 32-40, 42-50, 52, 53, and 64 are patentable under 35 U.S.C. § 103(a) over U.S. Patent No. 6,006,264 to Colby et. al. ("Colby"), in view of U.S. Patent No. 6,591,298 to Spicer et. al. ("Spicer").
- Whether Claims 9, 19, 41 and 51 are patentable under 35 U.S.C. §103(a) over Colby, in view of Spicer and further in view of U.S. Patent No. 6,526,283 to Jang ("Jang").

RECEIVED
CENTRAL FAX CENTER
DEC 05 2006

VII. Argument.

A. Rejection under 35 USC § 103(a) over Colby, in view of Spicer.

Claims 1-7, 12-17, 33-39, and 44-49

The Examiner rejected Claim 1 under 35 U.S.C. § 103(a) as unpatentable over Colby in view of Spicer. Claim 1 recites directing a first network client to one of the content servers if cost measurements are available that measure operational characteristics of the network based on communication between the first network client and one or more of the content servers, otherwise directing the first network client to one of the content servers using one or more cost measurements that measure operational characteristics of the network based on communication between a second client that is physically proximate to the first network client and one or more of the content servers.

In rejecting Claim 1, the Examiner alleged that Colby describes "directing the network client to a said one of said content servers based on communication between a client that is physically proximate to the network client and one or more of the plurality of content servers" and cited Figure 1 of Colby. Colby describes selecting a server based on the proximity of the requesting client to the servers. Column 3, lines 15-21 and Column 6, lines 42-43. Colby only considers the proximity of the requesting client and the servers. Colby does not consider the proximity of clients to one another. In contrast, Claim 1 requires the consideration of the proximity of a second client to a first network client to identity a server for the first network client. In particular, Claim 1 requires that the first network client is directed to one of the content servers based on cost measurements that measure operational characteristics of the network based on communication between a second client and one or more of the content servers where the second client is physically proximate to the first network client. Unlike Colby, the claimed invention does not select a content server based on the content server's physical proximity to the requesting client. Instead, the claimed invention selects a content server based on measurements associated with another client that is physically proximate to the requesting client.

The Examiner admitted that Colby does not describe a first network client and a second network client. However, the Examiner alleged that Spicer describes a first network client and a second network client at Figure 4, element 72, as well as Column 4, lines 30-55. The cited section of Spicer describes that multiple points of presence (POPs) can be connected to a network. Although Spicer describes that a given geographic location, such as a metropolitan area, is served by multiple network service providers and each network service provider provides a POP, the cited section of Spicer does not describe that measurements between one POP and a server can be used for routing communications from a separate POP that is physically proximate to the first POP.

The Examiner admitted that Colby does not teach cost measurements, but alleged that Spicer teaches cost measurements. Spicer discloses a method for measuring the performance of a web site by measuring download times at locations and times similar to that of users accessing the same web site. Column 1, lines 7-10 and lines 48-50. Although Spicer describes collecting performance measurements, it does not describe using information about one client to select a content server for another client. The cited section of Spicer describes multiple POPs within a given geographic location, such as a metropolitan area. Each POP is associated with a different network service provider. There is no suggestion in Spicer that performance data collected by one network service provider is shared with another network service provider. Thus, Spicer does not describe that performance measurements for one client are used to select a server for another client, where the clients are physically proximate.

The Examiner alleged that it would have been obvious to combine the teachings of Colby and Spicer because it would be useful to have "measurements to detect problems with content, network, a web server and a back-end system, or combinations thereof." It would not have been obvious to one of ordinary skill in the art at the time of the invention to combine the teachings of Colby and Spicer in the manner suggested by the Examiner. Colby describes selecting a server based on the server's physical or topological location, whereas Spicer describes selecting a server

based on web site performance measurements taken at specified intervals. There is no discussion in Colby or Spicer of how to resolve differences in the two methods, *i.e.* what to do if the performance measurement information contradicts the geographical or topological information.

Even if the teachings of Spicer and Colby are combined—despite the lack of teaching or motivation to do so—the claimed invention is not obtained. Combining the teachings of Spicer with the teachings of Colby would result in a hybrid method for directing a network client to requested content based on an evaluation of the performance of a web site measured by download times at various locations and times or based on the geographical or topological location of the server. The combination would not result in or suggest directing a first network client to one or more content servers using one or more cost measurements that measure operational characteristics of the network based on communications between a second client and the servers, where the second client is physically proximate to the first network client.

Claims 12, 33 and 44 also recite identifying a content server for a network client based on cost measurements that measure operational characteristics of the network based on communication between a client and one or more of the plurality of content servers, where the client is physically proximate to the network client. Claims 12, 33 and 44 are patentable over Colby and Spicer for the same reasons as Claim 1.

Claims 2-7 depend from Claim 1, Claims 13-17 depend from Claim 12, Claims 34-39 depend from Claim 33, and Claims 45-49 depend from Claim 44. The dependent claims are patentable over the cited references for at least the same reasons as the independent claims.

Claims 32 and 64

The Examiner rejected Claim 32 under 35 U.S.C. § 103(a) as unpatentable over Colby in view of Spicer. Claim 32 recites inferring operational characteristics associated with a plurality of network clients to an inferable network client by determining physical distances between said network clients and the inferable network client and computing a weighted average of said latency measurements as a function

of said distances, wherein said weighted average comprises an estimate of the latency between said network server and said inferable network client. The Examiner alleged that Colby describes determining physical distances between said network clients and the inferable network client. The sections of Colby relied upon by the Examiner describe the selection of a server and the determination of flow pipe bandwidth. Colby describes routing based, in part, on the proximity of the server to the requesting client. As discussed above in connection with Claim 1, Colby does not describe using measurements associated with one client to infer or estimate performance for a different or inferable client. The Examiner did not cite any sections of Spicer in rejecting Claim 32. As discussed above in connection with Claim 1, Spicer does not describe using measurements associated with one client to route another client to a server. Thus, Claim 32 is patentable over Colby and Spicer. Claim 64 also recites inferring operational characteristics associated with a plurality of network clients to an inferable network client by determining physical distances between said network clients and the inferable network client and computing a weighted average of said latency measurements to estimate the latency between said network server and the inferable network client. Claim 64 is patentable over Colby and Spicer for the same reasons as Claim 32.

Claims 10, 20, 42, and 52

The Examiner rejected Claim 10 under 35 U.S.C. § 103(a) as unpatentable over Colby in view of Spicer. Claim 10 recites inferring network performance by determining a weighted average of network performance between said content server and other network clients based on the physical proximity of the other network clients to said first network client. The sections of Colby relied upon by the Examiner in rejecting Claim 10 describe the selection of a server and the determination of flow pipe bandwidth. Colby describes routing based, in part, on the proximity of the server to the requesting client. As discussed above in connection with Claim 1, Colby does not describe using measurements associated with one client to infer or estimate performance for a different client. Claims 20, 42 and 52 also recite inferring network

performance of the first network client based on a weighted average of network performance between the content server and other network clients, where the other network clients are physically proximate to the first network client and are patentable for the same reasons as Claim 10.

Claim 10 depends from Claim 1, Claims 20 depends from Claim 12, Claim 42 depends from Claim 33, and Claim 52 depends from Claim 44. The dependent claims are patentable over the cited references for at least the same reasons as the independent claims.

Claims 11, 21, 43, and 53

The Examiner rejected Claim 11 under 35 U.S.C. § 103(a) as unpatentable over Colby in view of Spicer. Claim 11 recites determining physical distances between said network clients and the first network client and computing a weighted average of said latency measurements as a function of said distances, wherein said weighted average comprises an estimate of the latency between said network server and said first network client. The sections of Colby relied upon by the Examiner to reject Claim 11 describe the selection of a server and the determination of flow pipe bandwidth. Colby describes routing based, in part, on the proximity of the server to the requesting client. As discussed above in connection with Claim 1, Colby does not describe using measurements associated with one client to infer or estimate performance for a different client. Claims 21, 43, and 53 also recite determining physical distances between said network clients and the first network client and computing a weighted average of said latency measurements as a function of said distances, wherein said weighted average comprises an estimate of the latency between said network server and said first network client and are patentable for the same reasons as Claim 11.

Claim 11 depends from Claim 1, Claim 21 depends from Claim 12, Claim 43 depends from Claim 33, and Claim 53 depends from Claim 44. The dependent claims are patentable over the cited references for at least the same reasons as the independent claims.

Claims 8, 18, 40 and 50.

The Examiner rejected Claim 8 under 35 U.S.C. § 103(a) as unpatentable over Colby in view of Spicer. Claim 8 recites determining the location of said first network client by circular intersection. The Examiner rejected Claim 8 alleging that Figure 1A of Colby describes circular intersection. Figure 1A of Colby illustrates a conventional IP network where servers are connected to routers at the edges of the network and each router is connected to one or more routers. Figure 1A of Colby does not describe circular intersection, and thus, does not describe the invention of Claim 8. Claims 18, 40 and 50 also recite determining the location of the first network client by circular intersection and are patentable for the same reasons as Claim 8.

Claim 8 depends from Claim 1, Claim 18 depends from Claim 12, Claim 40 depends from Claim 33 and Claim 50 depends from claim 44. The remarks made above to distinguish the independent claims from Colby and Spicer are equally applicable to distinguish the dependent claims from the cited references.

B. Rejection under 35 USC § 103 (a) over Colby, Spicer and Jang.**Claims 9, 19, 41 and 51**

The Examiner rejected Claim 9 under 35 USC § 103 (a) as unpatentable over Colby in view of Spicer and Jang.² The Examiner alleged that it would have been obvious to combine Colby, Spicer and Jang to provide a plurality of intersecting circles using distance equivalents as the radii of the circles with the network server locations as the centers and determining the physical location of the client from the intersection of the circles "because it would have an efficient system that can provide specific degree or amount of separation between two points, lines, surfaces or objects or an advance along a route measured linearly."

² It is believed that the Examiner relied upon Colby, Spicer and Jang to reject Claim 9. The language of the Office action is inconsistent since at one point the Examiner cites only Colby and Jang, whereas at other points the Examiner cites Colby, Spicer and Jang. Since Claim 9 depends from Claim 1 and the Examiner relied upon Colby and Spicer in rejecting Claim 1, it is believed that the rejection of Claim 9 is based on Colby, Spicer and Jang.

It would not have been obvious to one of ordinary skill in the art at the time of the invention to combine the teachings of Colby, Spicer and Jang in the manner suggested by the Examiner. There is no suggestion that the inventions described in Colby, Spicer or Jang would benefit from a system that provides specific degrees of separation between two points, lines, surfaces or objects or an advance along a route, as alleged by the Examiner. Jang is directed to determining the location of a mobile phone, whereas Colby and Spicer are directed to routing in computer networks so that the references are not all in the same technical field. The signal formats and protocols used on the computer networks described by Colby and Spicer are different than those used in a mobile phone network, such as that described by Jang. Although Colby considers the proximity of the server to the requesting client, Colby describes determining whether the server and client reside on the same continent. *See e.g.* Column 17, lines 59-66. As described in more detail in connection with Claim 1, Spicer uses performance measurements to determine routing and has no need to determine the physical location of the devices on the network. Thus, neither Colby nor Spicer suggest the need to use circular intersection.

Even if the teachings of Spicer, Colby, and Jang are combined—despite the lack of teaching or motivation to do so—the claimed invention is not obtained. The combination would not result in or suggest directing a first network client to one or more content servers using one or more cost measurements that measure operational characteristics of the network based on communications between a second client and the content servers, where the second client is physically proximate to the first network client, and where the location of the first network client is determined using circular intersection. Claims 19, 41, and 51 also recite determining the location of the first network client by circular intersection and are patentable for the same reasons as Claim 9.

Claim 9 depends from Claim 1, Claim 19 depends from Claim 12, Claim 41 depends from Claim 33 and Claim 51 depends from Claim 44.. The remarks made above to distinguish the independent claims from Colby and Spicer are equally applicable to distinguish the dependent claims from the cited references.

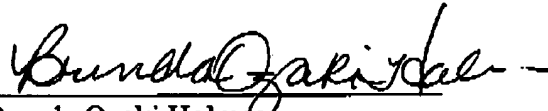
Conclusion

For at least the reasons given above, it is respectfully requested that the final rejection be reversed and Claims 1-21, 32-53 and 64 be allowed.

Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account 11-0855, and please credit any excess fees to such deposit account.

Respectfully submitted,

By:


Brenda Ozaki Holmes
Registration No. 40,339

Kilpatrick Stockton LLP
1100 Peachtree Street
Suite 2800
Atlanta, Georgia 30309
(404) 815-6500

RECEIVED
CENTRAL FAX CENTER
DEC 05 2006

VIII. Claims Appendix.

1. (Previously Presented) A method for directing a first network client requesting access to content to one of a plurality of content servers that can provide said content, comprising:

if one or more cost measurements are available that measure operational characteristics of the network based on communication between the first network client and one or more of the plurality of content servers, then directing the first network client to a said one of said content servers based on the one or more cost measurements;

otherwise, directing the first network client to a said one of said content servers using one or more cost measurements that measure operational characteristics of the network based on communication between a second client that is physically proximate to the first network client and one or more of the plurality of content servers.

2. (Previously Presented) A method as recited in claim 1, further comprising:

obtaining a new cost measurement when said first network client accesses said content server; and

using said new cost measurement as an indicator of operational characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

3. (Original) A method as recited in claim 1, wherein said content servers are associated with a network server having an identity, and wherein said network client requests content from said network server, and further comprising:

mapping the identity of the network server to said content servers.

4. (Previously Presented) A method as recited in claim 1, further comprising measuring network performance between said first network client and a said one of said content servers.

5. (Original) A method as recited in claim 1, wherein an attribute of network performance comprises network latency.

6. (Previously Presented) A method as recited in claim 5, wherein network latency is measured passively by determining the time between a syn ack message sent by said first network client and an ack message sent by one of said content servers.

7. (Previously Presented) A method as recited in claim 4, further comprising measuring network performance between said first network client and another of said content servers.

8. (Previously Presented) A method as recited in claim 1, further comprising determining the location of said first network client by circular intersection.

9. (Previously Presented) A method as recited in claim 8, wherein said circular intersection comprises:

- (a) measuring the time that it takes for data to move from a plurality of network server locations to said first client;
- (b) converting said times to distance equivalents;
- (c) determining a plurality of intersecting circles, wherein said distance equivalents are used as the radii of the circles and said network server locations are used as the centers of the circles; and
- (d) determining the physical location of said first network client from the intersection of said circles.

10. (Previously Presented) A method as recited in claim 1, further comprising inferring network performance of serving said first network client from said content server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said network client and performance of said content server for said other network clients.

11. (Previously Presented) A method as recited in claim 1, further comprising:

(a) measuring network latency between a content server and a plurality of other network clients;

(b) determining physical distances between said other network clients and said first network client;

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said content server and said first network client; and

(d) inferring operational characteristics associated with a plurality of network clients to said first network client using said weighted average.

12. (Previously Presented) A method for directing a network client requesting access to content from a network server to one of a plurality of content servers that can provide said content, each said content server having an address, said network server having an identity, said method comprising:

(a) identifying the network server associated with the content requested by said network client;

(b) if one or more cost measurements are available that measure operational characteristics of the network based on communication between the first network client and one or more of the plurality of content servers, then identifying a said one of said content servers based on said identity of said network server and the one or more cost measurements; otherwise, identifying a said one of said content servers based on

said identity of said network server and one or more cost measurements that measure operational characteristics of the network based on communication between a second client that is physically proximate to the first network client and one or more of the plurality of content servers; and

(c) providing the first network client with the address of said content server identified in step (b).

13. (Previously Presented) A method as recited in claim 12, further comprising:

(d) obtaining a new cost measurement when said first network client accesses said content server; and

(e) using said new cost measurement as an indicator of operational characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

14. (Previously Presented) A method as recited in claim 12, further comprising measuring network performance between said first network client and a said one of said content servers.

15. (Original) A method as recited in claim 12, wherein an attribute of network performance comprises network latency.

16. (Previously Presented) A method as recited in claim 15, wherein network latency is measured passively by determining the time between a syn ack message sent by said first network client and an ack message sent by said one of said content servers.

17. (Previously Presented) A method as recited in claim 14, further comprising measuring network performance between said first network client and another of said content servers.

18. (Previously Presented) A method as recited in claim 12, determining the location of said first network client by circular intersection.

19. (Previously Presented) A method as recited in claim 18, wherein said circular intersection comprises:

- (a) measuring the time that it takes for data to move from a plurality of network server locations to said first network client;
- (b) converting said times to distance equivalents;
- (c) determining a plurality of intersecting circles, wherein said distance equivalents are used as the radii of the circles and said network server locations are used as the centers of the circles; and
- (d) determining the physical location of said first network client from the intersection of said circles.

20. (Previously Presented) A method as recited in claim 12, further comprising inferring network performance of serving said first network client from said content server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said first network client and performance of said content server for said other network clients.

21. (Previously Presented) A method as recited in claim 12, further comprising:

- (a) measuring network latency between a content server and a plurality of other network clients;
- (b) determining physical distances between said other network clients and said first network client;
- (c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said content server and said first network client; and
- (d) inferring operational characteristics associated with said other network clients to said network client using said weighted average.

22. – 31. (Cancelled)

32. (Previously Presented) A method for inferring operational characteristics associated with a plurality of network clients to an inferable network client, comprising:

(a) measuring network latency between a network server and the plurality of network clients;

(b) determining physical distances between said network clients and the inferable network client; and

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said network server and said inferable network client.

33. (Previously Presented) A system for directing a first network client requesting access to content to one of a plurality of content servers that can provide said content, comprising:

a programmed data processor; and

programming associated with said programmed data processor for:

determining whether one or more cost measurements are available that measure operational characteristics of the network based on communication between the first network client and one or more of the plurality of content servers,

if so, then directing the first network client to a said one of said content servers based on the one or more cost measurements;

if not, then directing the first network client to a said one of said content servers using one or more cost measurements that measure operational characteristics

of the network based on communication between a second client that is physically proximate to the first network client and one or more of the plurality of content servers.

34. (Previously Presented) A system as recited in claim 33, further comprising programming associated with said programmed data processor for:

obtaining a new cost measurement when said first network client accesses said content server; and

using said new cost measurement as an indicator of operational characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

35. (Previously Presented) A system as recited in claim 33, wherein said content servers are associated with a network server having an identity, and wherein said first network client requests content from said network server, and further comprising:

programming associated with said programmed data processor mapping the identity of the network server to said content servers.

36. (Previously Presented) A system as recited in claim 33, further comprising programming associated with said programmed data processor for measuring network performance between said first network client and a said one of said content servers.

37. (Original) A system as recited in claim 33, wherein an attribute of network performance comprises network latency.

38. (Previously Presented) A system as recited in claim 37, wherein network latency is measured passively by determining the time between a syn ack message sent by said first network client and an ack message sent by one of said content servers.

39. (Previously Presented) A system as recited in claim 36, further comprising programming associated with said programmed data processor for measuring network performance between said first network client and another of said content servers.

40. (Previously Presented) A system as recited in claim 33, further comprising programming associated with said programmed data processor for determining the location of said first network client by circular intersection.

41. (Previously Presented) A system as recited in claim 40, wherein said circular intersection comprises:

(a) measuring the time that it takes for data to move from a plurality of network server locations to said first network client;

(b) converting said times to distance equivalents;

(c) determining a plurality of intersecting circles, wherein said distance equivalents are used as the radii of the circles and said network server locations are used as the centers of the circles; and

(d) determining the physical location of said network client from the intersection of said circles.

42. (Previously Presented) A system as recited in claim 33, further comprising programming associated with said programmed data processor for inferring network performance of serving said first network client from said content server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said first network client and performance of said content server for said other network clients.

43. (Previously Presented) A system as recited in claim 33, further comprising programming associated with said programmed data processor for

(a) measuring network latency between a content server and a plurality of other network clients;

(b) determining physical distances between said other network clients and said first network client;

(c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said content server and said first network client; and

(d) inferring operational characteristics associated with a plurality of network clients to said first network client using said weighted average.

44. (Previously Presented) A system for directing a first network client requesting access to content from a network server to one of a plurality of content servers that can provide said content, each said content server having an address, said network server having an identity, said method comprising:

- (a) a programmed data processor; and
- (b) programming associated with said programmed data processor for
 - (i) identifying the network server associated with the content requested by said first network client;
 - (ii) if one or more cost measurements are available that measure operational characteristics of the network based on communication between the network client and one or more of the plurality of content servers, then identifying a said one of said content servers based on said identity of said network server and the one or more cost measurements; and
 - (iii) otherwise, identifying a said one of said content servers based on said identity of said network server and one or more cost measurements that measure operational characteristics of the network based on communication between a client that is physically proximate to the first network client and one or more of the plurality of content servers; and
- (c) providing the first network client with the address of said content server identified in step (b).

45. (Previously Presented) A system as recited in claim 44, further comprising programming associated with said programmed data processor for:

obtaining a new cost measurement when said first network client accesses said content server; and

using said new cost measurement as an indicator of operational characteristics of the network in connection with subsequent requests for access to said content that can be provided by said content server.

46. (Previously Presented) A system as recited in claim 44, further comprising programming associated with said programmed data processor for measuring network performance between said first network client and a said one of said content servers.

47. (Original) A system as recited in claim 46, wherein an attribute of network performance comprises network latency.

48. (Previously Presented) A system as recited in claim 47, wherein network latency is measured passively by determining the time between a syn ack message sent by said first network client and an ack message sent by said one of said content servers.

49. (Previously Presented) A system as recited in claim 46, further comprising programming associated with said programmed data processor for measuring network performance between said first network client and another of said content servers.

50. (Previously Presented) A system as recited in claim 44, further comprising programming associated with said programmed data processor for determining the location of said network client by circular intersection.

51. (Previously Presented) A system as recited in claim 50, wherein said circular intersection comprises:

- (a) measuring the time that it takes for data to move from a plurality of network server locations to said first network client;
- (b) converting said times to distance equivalents;
- (c) determining a plurality of intersecting circles, wherein said distance equivalents are used as the radii of the circles and said network server locations are used as the center of the circles; and
- (d) determining the physical location of said network client from the intersection of said circles.

52. (Previously Presented) A system as recited in claim 44, further comprising programming associated with said programmed data processor for inferring network performance of serving said first network client from said content

server by determining a weighted average of network performance between said content server and other network clients based on physical proximity of said other network clients to said first network client and performance of said content server for said other network clients.

53. (Previously Presented) A system as recited in claim 44, further comprising programming associated with said programmed data processor for:

- (a) measuring network latency between a content server and a plurality of other network clients;
- (b) determining physical distances between said other network clients and said first network client;
- (c) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said content server and said network client; and
- (d) inferring operational characteristics associated with a plurality of network clients to said network client using said weighted average.

54. – 63. (Cancelled)

64. (Previously Presented) A system for inferring operational characteristics associated with a plurality of network clients to an inferable network client, comprising:

- (a) a programmed data processor; and

- (b) programming associated with said programmed data processor for
 - (i) measuring network latency between a network server and the plurality of network clients;
 - (ii) determining physical distances between said network clients and the inferable network client; and
 - (iii) computing a weighted average of said latency measurements as a function of said distances, wherein said weighed average comprises an estimate of the latency between said network server and said inferable network client.

RECEIVED
CENTRAL FAX CENTER
DEC 05 2006

IX. Evidence Appendix.

None.

RECEIVED
CENTRAL FAX CENTER
DEC 05 2006

X. Related Proceedings Appendix.

None.